High Performance Computing in Science and Engineering: the Tree and the Fruit

David Keyes

Dept. of Applied Physics & Applied Mathematics, Columbia University

Div. of Mathematical and Computer Sciences & Engineering, KAUST

Simulation: third pillar of scientific discovery

- Theory
- Experiment
- Simulation

"theoretical experiments"

Computational simulation :

"a means of scientific discovery that employs a computer system to simulate a physical system according to laws derived from theory and experiment"

not an "other" but a hybrid and a platform for integration of both

Simulation driven by price and capability

By the Gordon Bell Prize, simulation *cost per performance* has improved by nearly a million times in two decades. Performance on *real applications* (e.g., mechanics, materials, petroleum reservoirs, gravitation) has improved *more* than a million times.

Gordon Bell Prize: Price Performance **Year**	**Cost per delivered Gigaflop/s**
1989	$2,500,000
1999	$6,900
2009	$8

Gordon Bell Prize: Peak Performance **Year**	**Gigaflop/s delivered to applications**
1988	1
1998	1,020
2008	1,350,000

Whimsical remarks on simulation progress measured by Bell Prizes, since 1988

- **If similar improvements in *speed* (10^6) had been realized in the airline industry, a 15-hour flight (e.g., JFK-NRT) would require one-twentieth of a second today**

- **If similar improvements in *storage* (10^4) had been realized in the publishing industry, our office bookcases could hold the book portion of the collection of the U.S. Library of Congress (~22M volumes)**

- **If similar reductions in *cost* ($>10^5$) had been realized in the higher education industry, tuition room and board (at a college in the USA) would cost about $0.20 per year**

Thought experiment:
How to use peanuts as price per ton falls?

- **In 2012, at $1,150./ton:**
  - **make sandwiches**
- **By 2015, at $115./ton:**
  - **make recipe substitutions**
- **By 2018, at $11.50/ton:**
  - **use as feedstock for plastics, etc.**
- **By 2021, at $1.15/ton:**
  - **heat homes**
- **By 2024, at $0.115/ton:**
  - **pave roads ☺**

The cost of computing has been on a curve like this for two decades and promises to continue. Like everyone else, scientists and engineers plan increasing uses for it…

1979: Computational Fluid Dynamics for B767

Much CFD penetration. Opportunities exist for higher accuracy and expanded complexity

Some CFD penetration. Opportunities exist for large increases in design process speed and application

CFD penetration opportunity

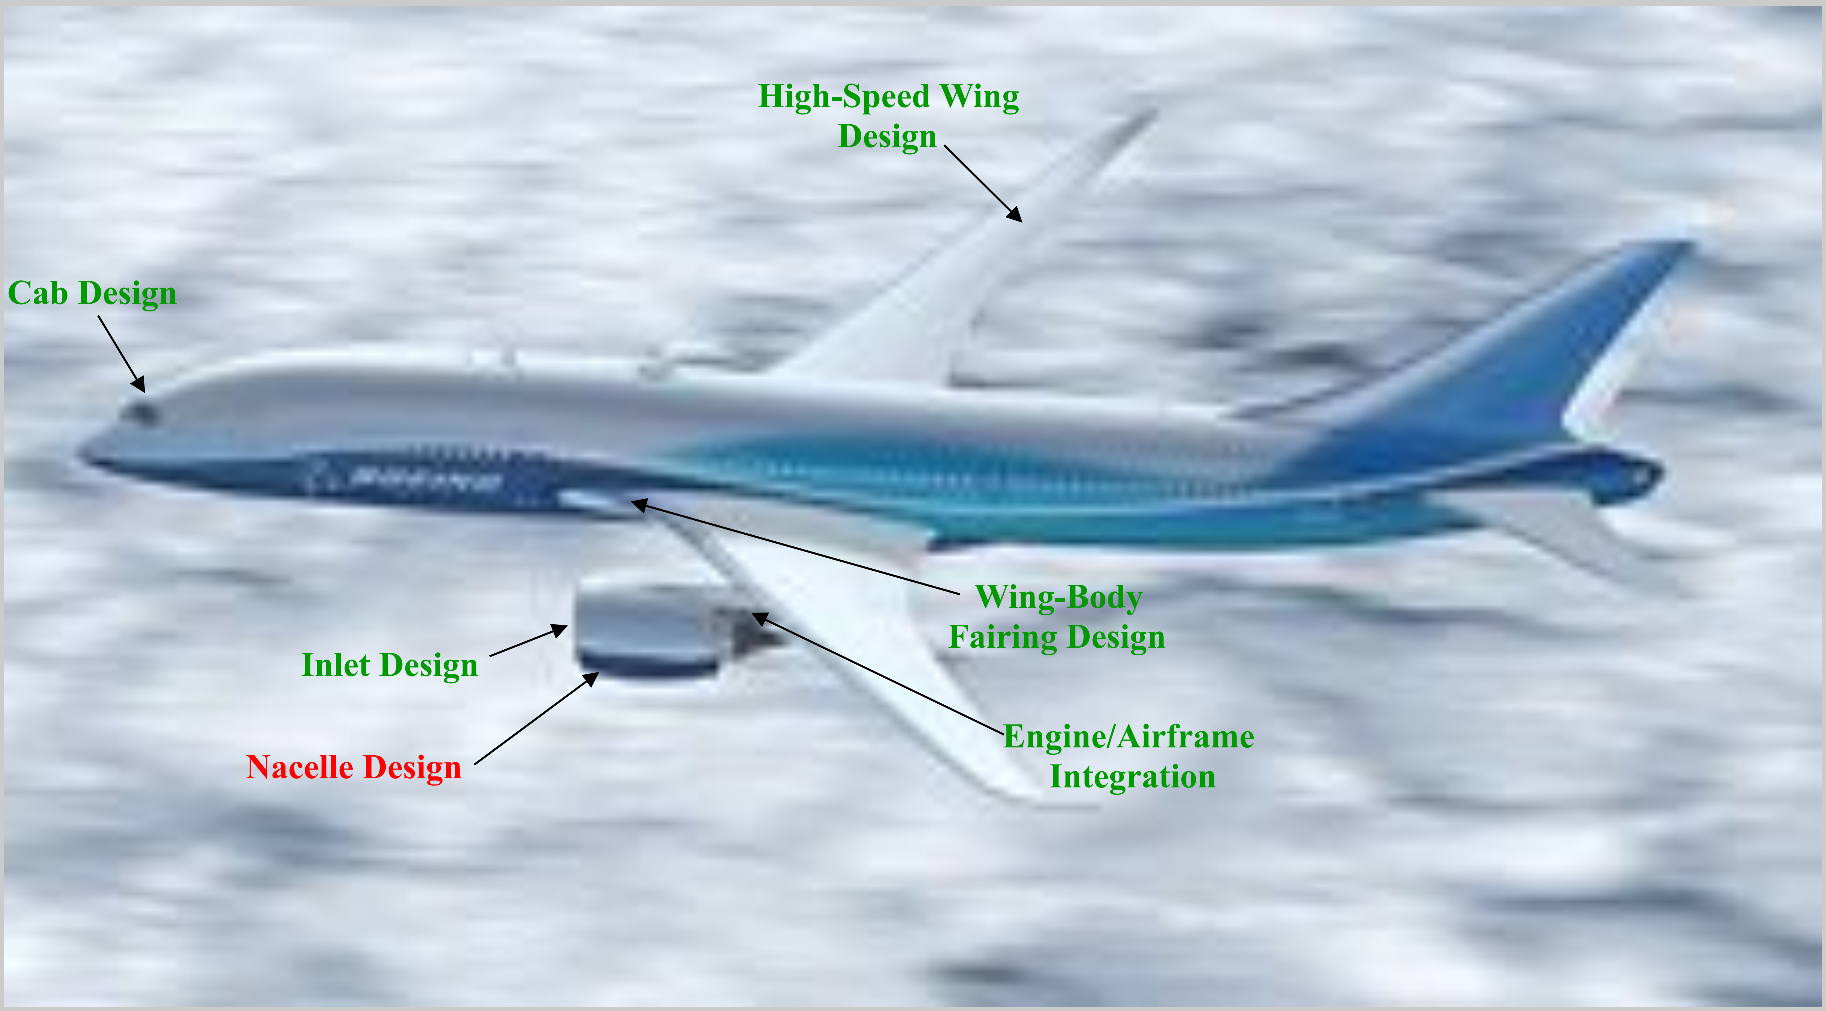

c/o Douglas Ball, Boeing

2005: Computational Fluid Dynamics for B787

Much CFD penetration. Opportunities exist for higher accuracy and expanded complexity

Some CFD penetration. Opportunities exist for large increases in design process speed and application

CFD penetration opportunity

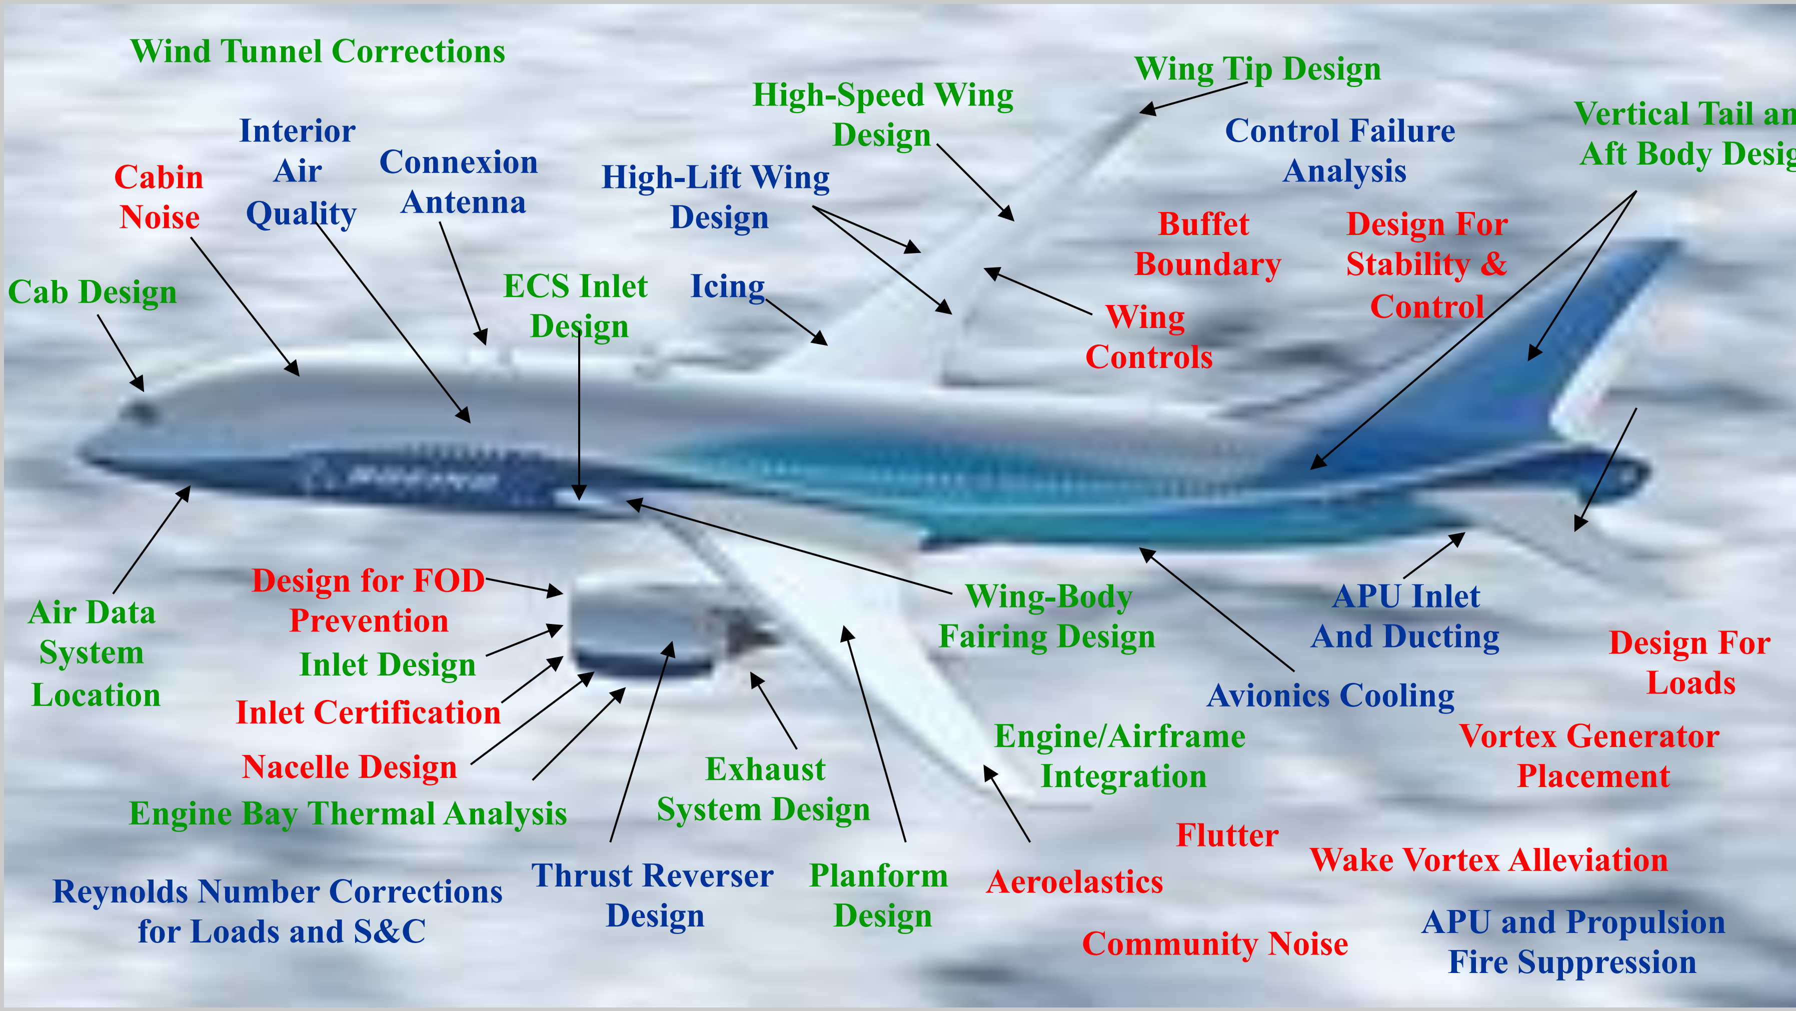

c/o Douglas Ball, Boeing

2011 buyer driving factors in HPC

Top Reasons for Acquiring HPC Systems	
Average rankings, 10 -- most important, 1 -- least important	
Attribute:	**Ranking**
Ability to do new/ better science	8.18
Ability to run larger problems	7.69
Performance on OUR applications	7.38
Throughput	7.17
Price/ Performance	6.99
To improve our competitiveness	6.06
Total cost of ownership	6.03
Capacity mgmt	5.29
Regulations/ certification	3.08
Source: IDC, 2011	

The imperative of simulation

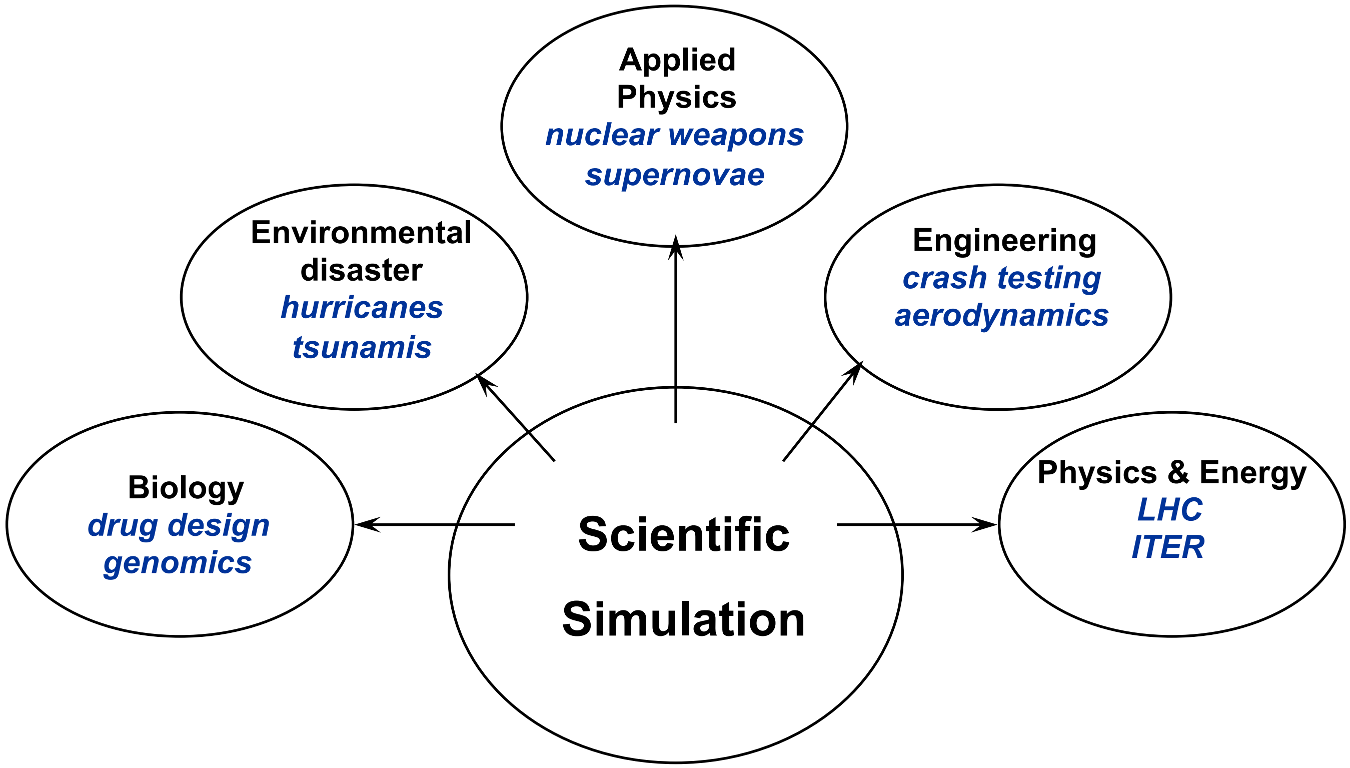

In these, and many other areas, simulation is an important complement to experiment.

The imperative of simulation

Experiments controversial

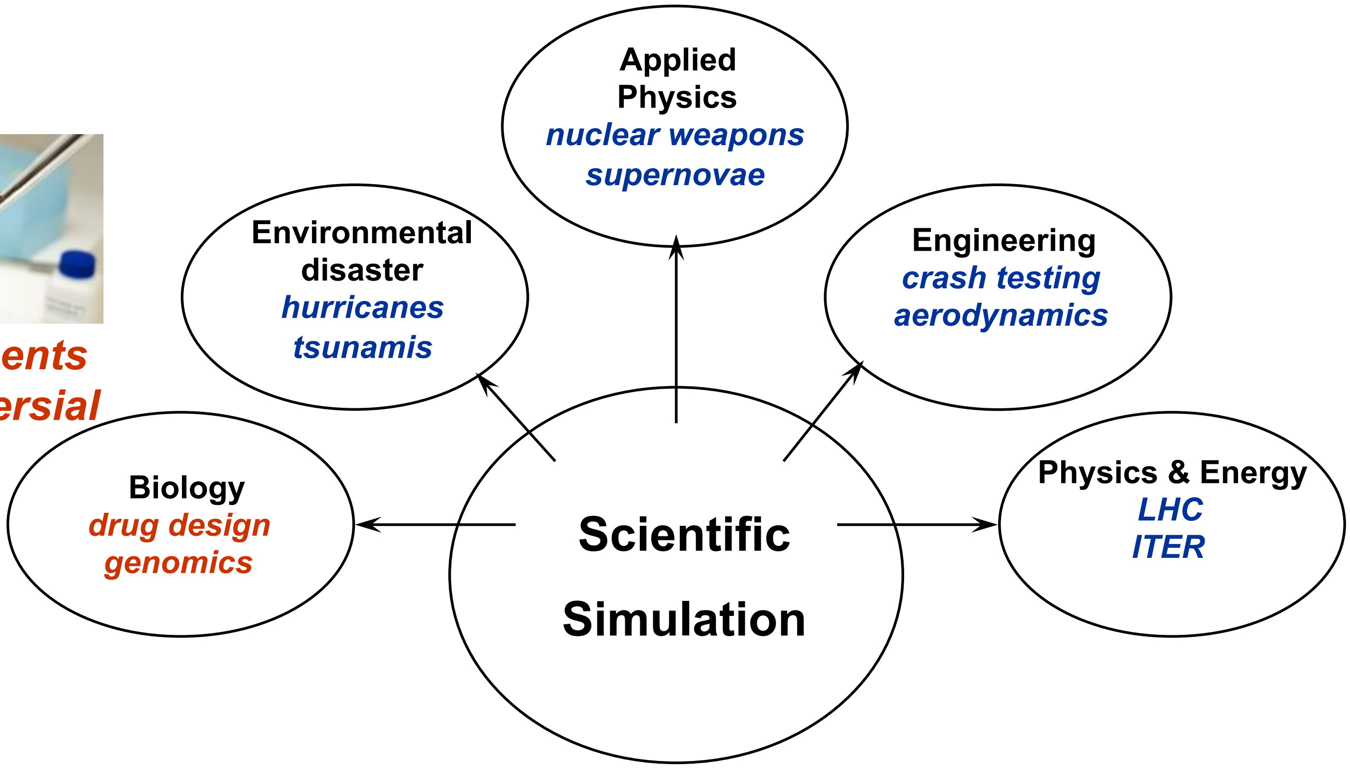

In these, and many other areas, simulation is an important complement to experiment.

The imperative of simulation

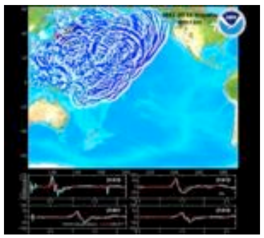

“Experiments” dangerous

Experiments controversial

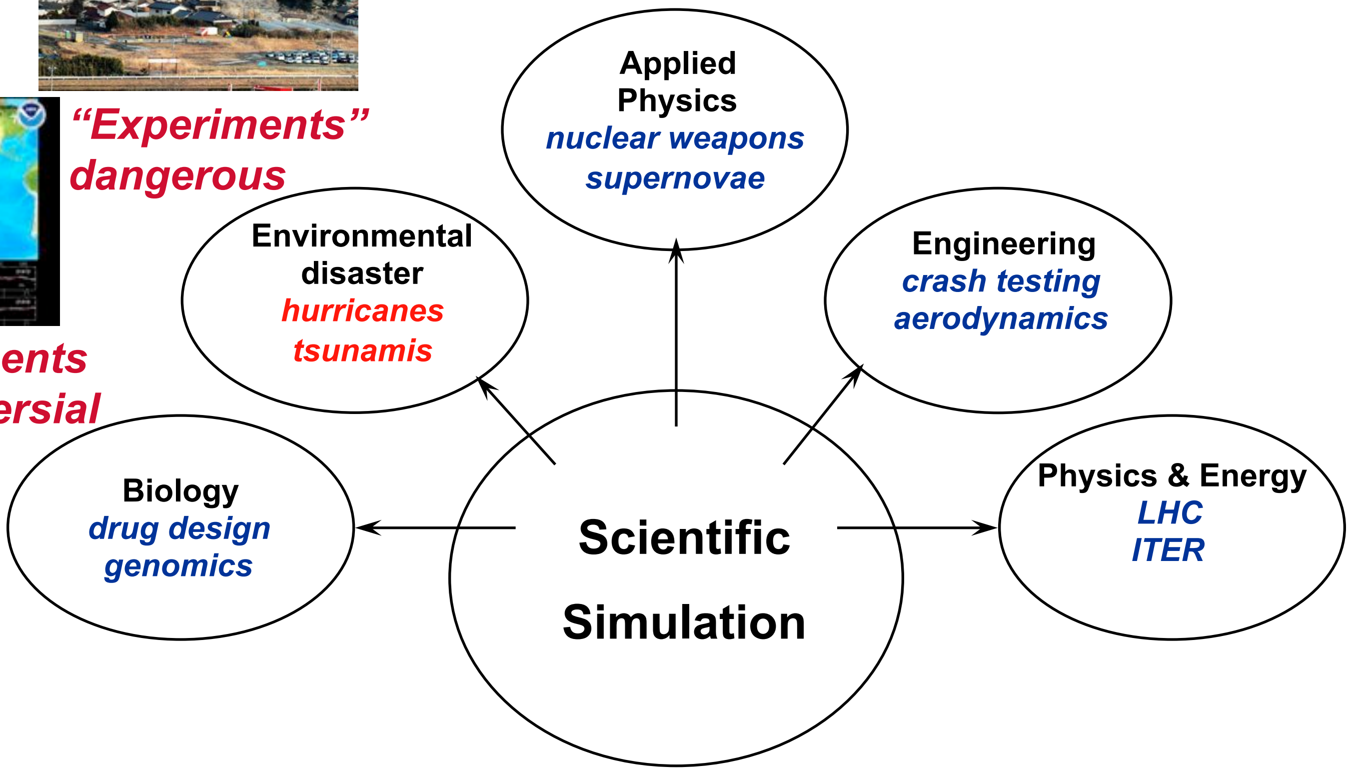

In these, and many other areas, simulation is an important complement to experiment.

The imperative of simulation

Experiments prohibited or impossible

"Experiments" dangerous

Experiments controversial

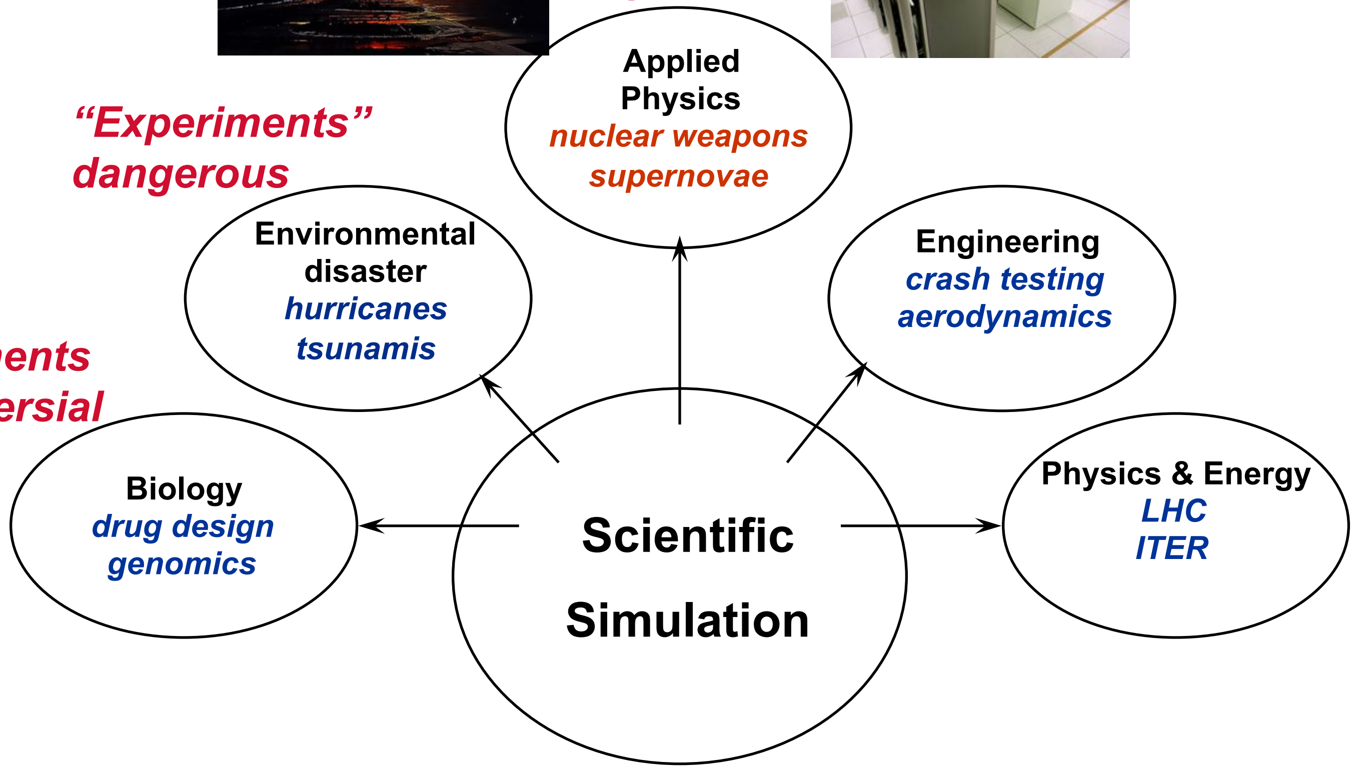

In these, and many other areas, simulation is an important complement to experiment.

The imperative of simulation

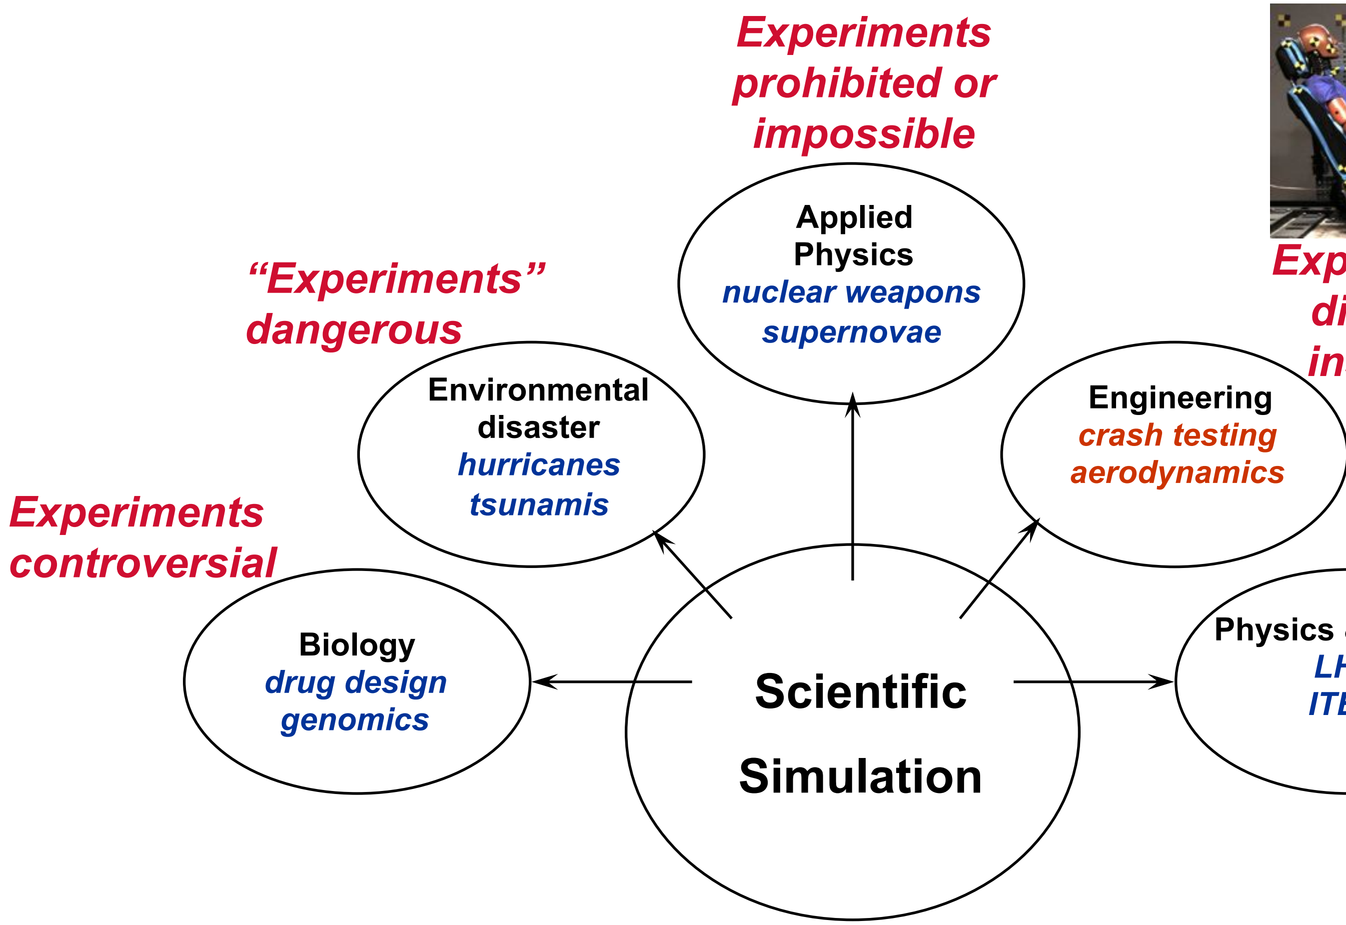

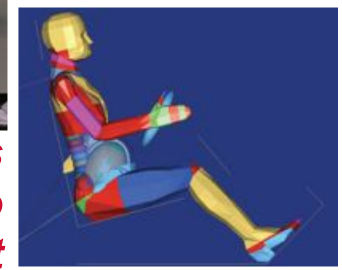

In these, and many other areas, simulation is an important complement to experiment.

The imperative of simulation

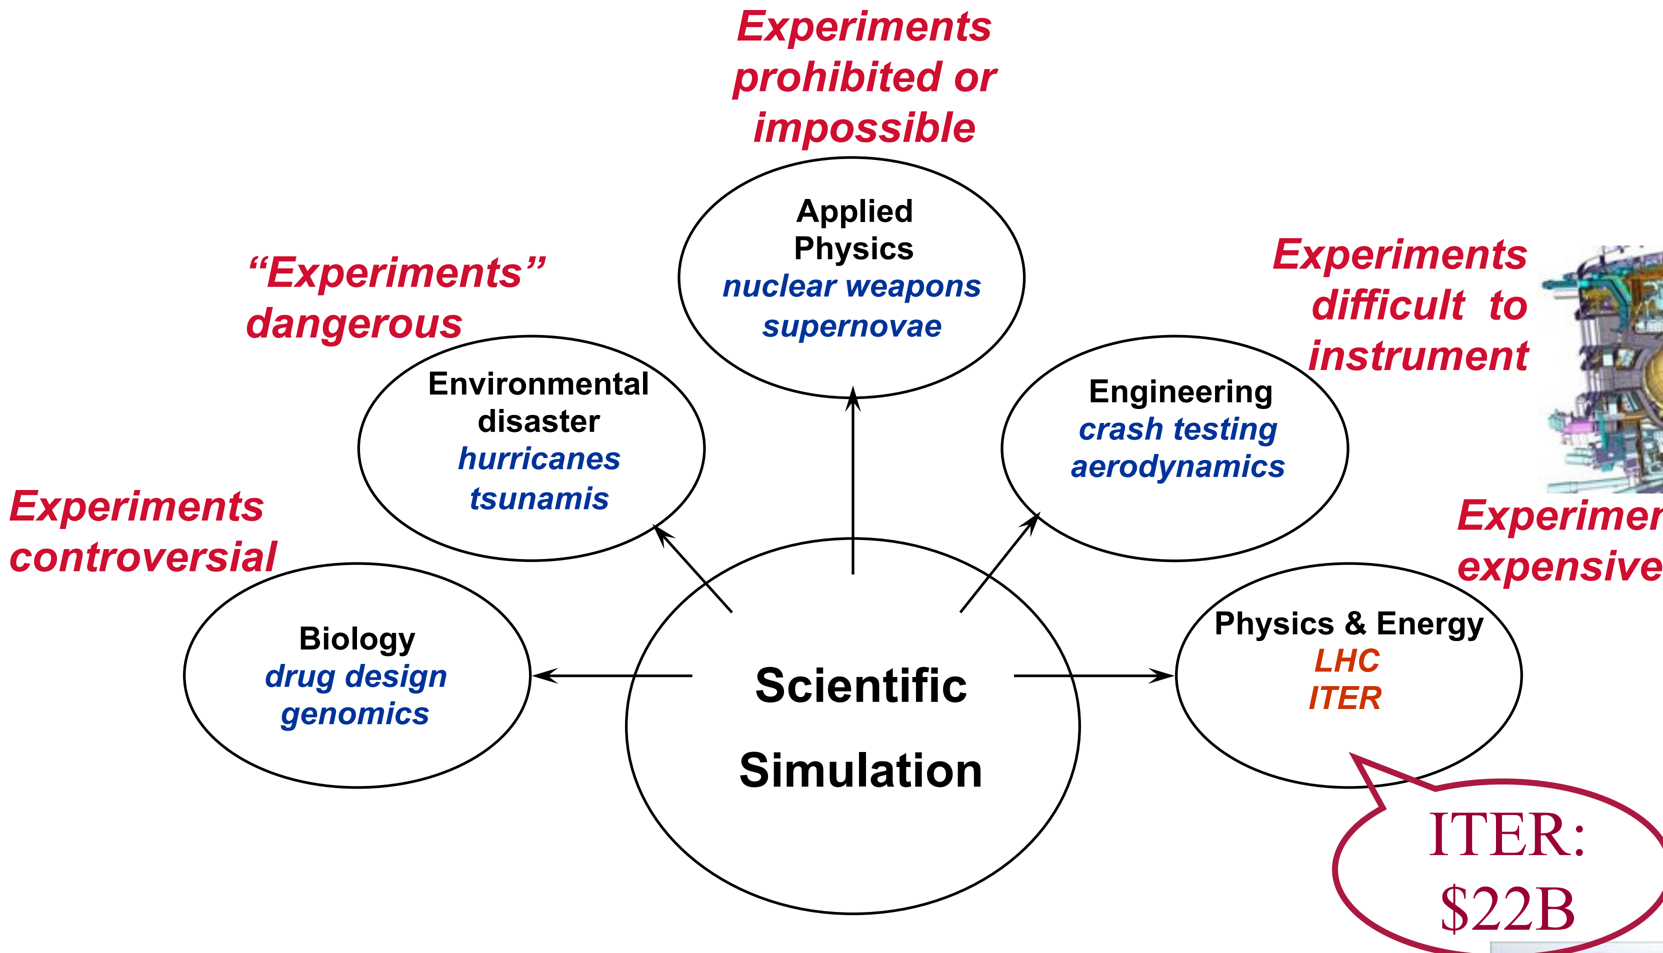

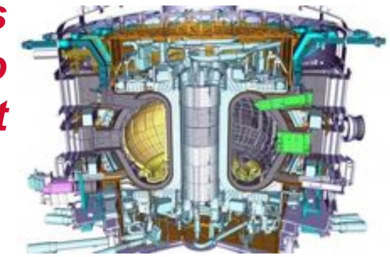

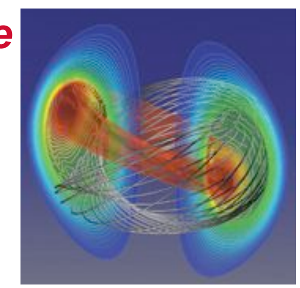

In these, and many other areas, simulation is an important complement to experiment.

Balance shift in modality of scientific discovery

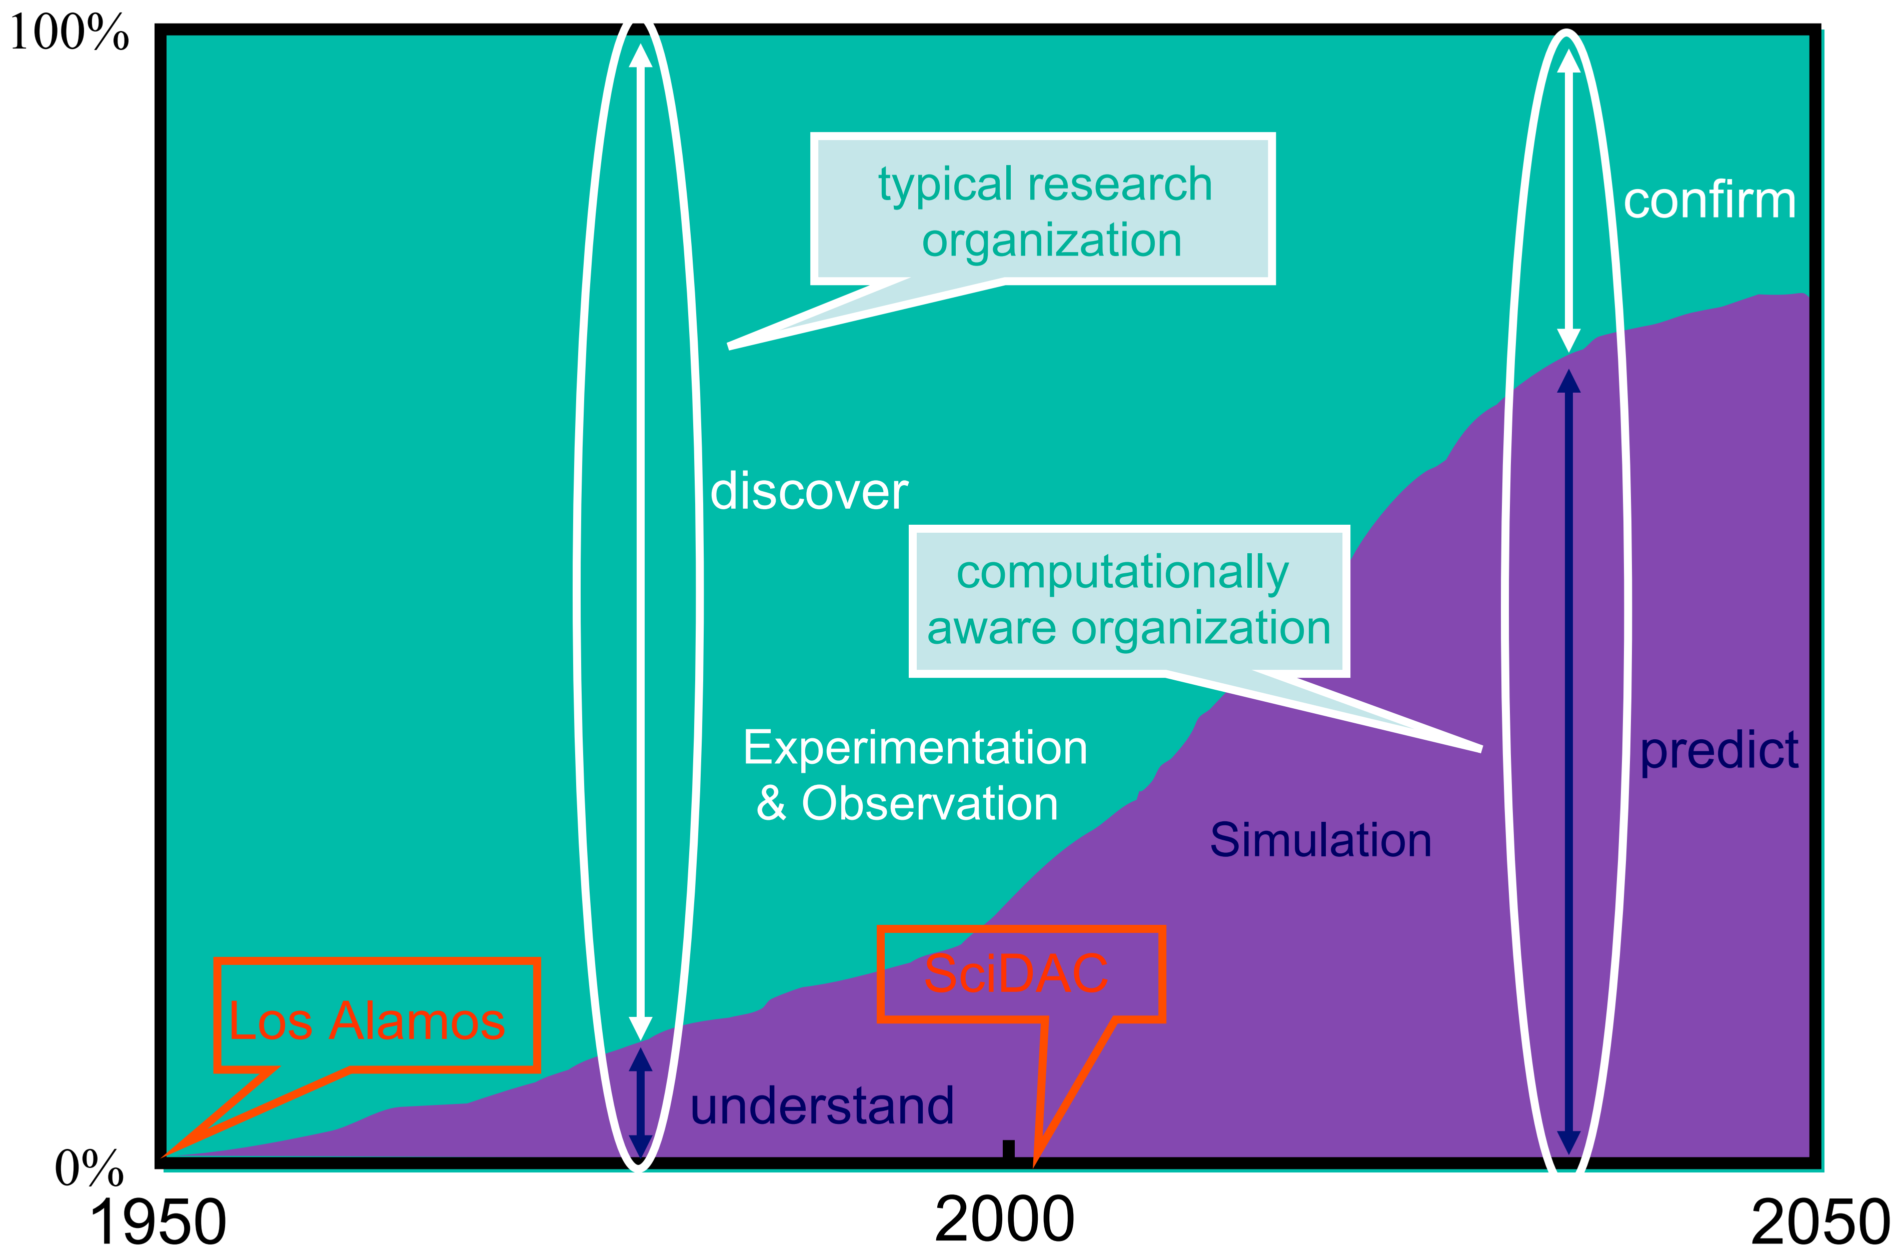

Moore's Law: exponential growth in time

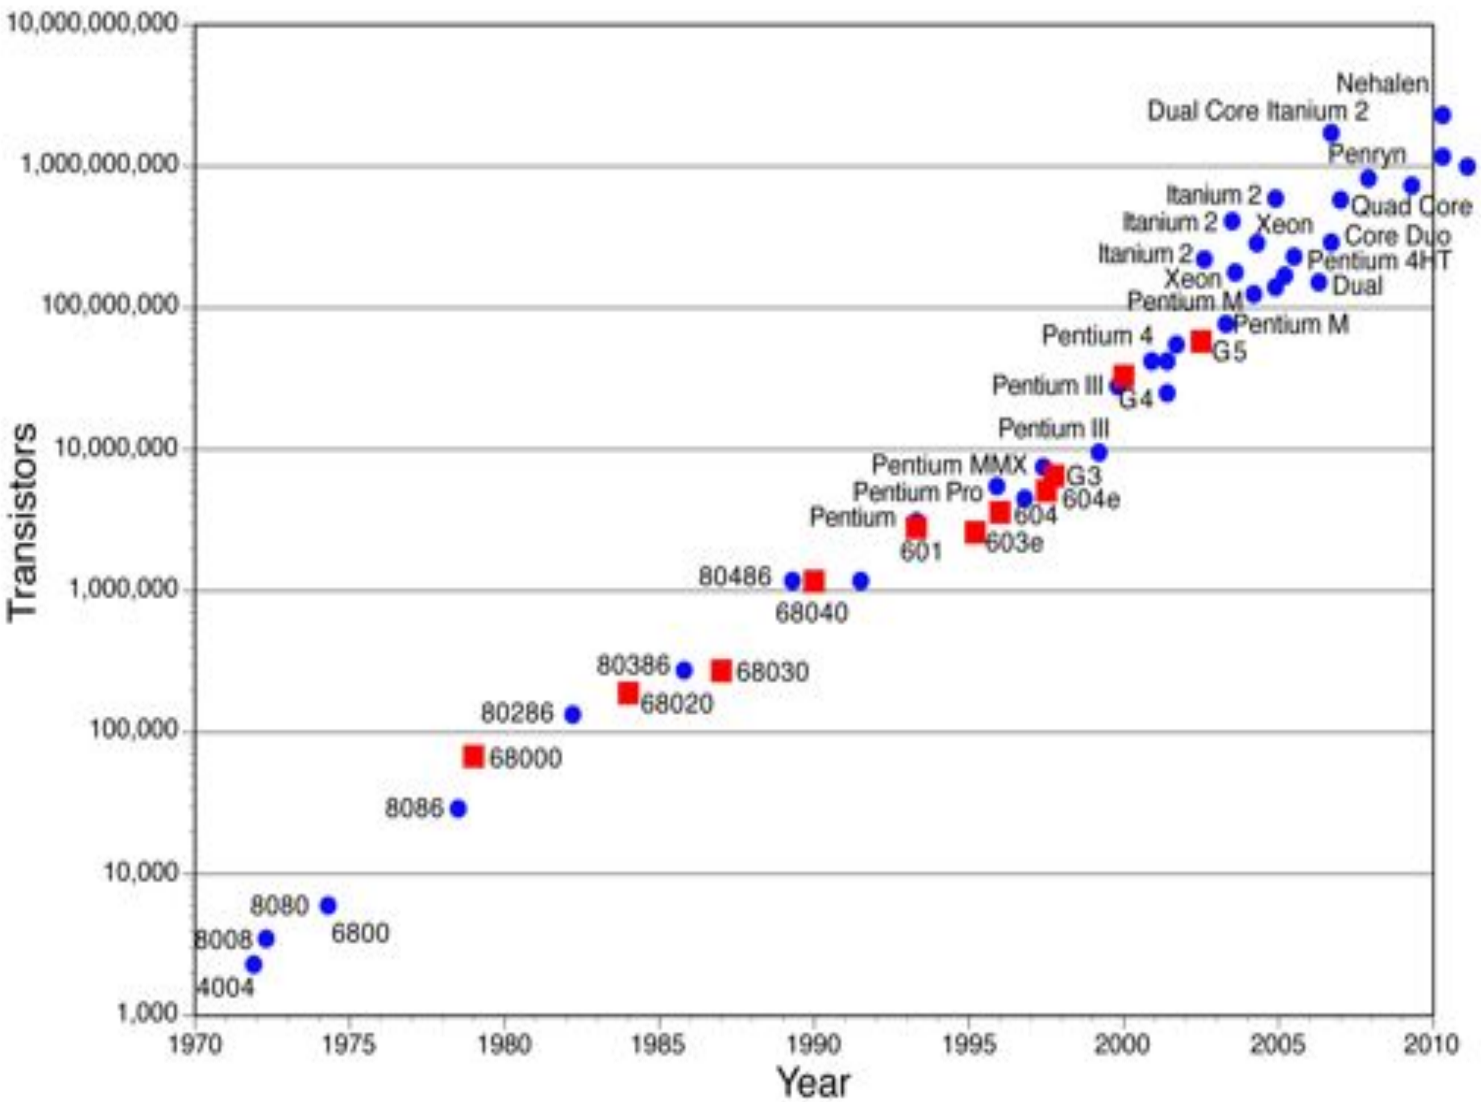

Attributed to Gordon Moore of Intel from a paper in 1965 projecting CMOS transistor density, the term is applied today throughout science and technology

"Moore's Law" for fusion energy simulations

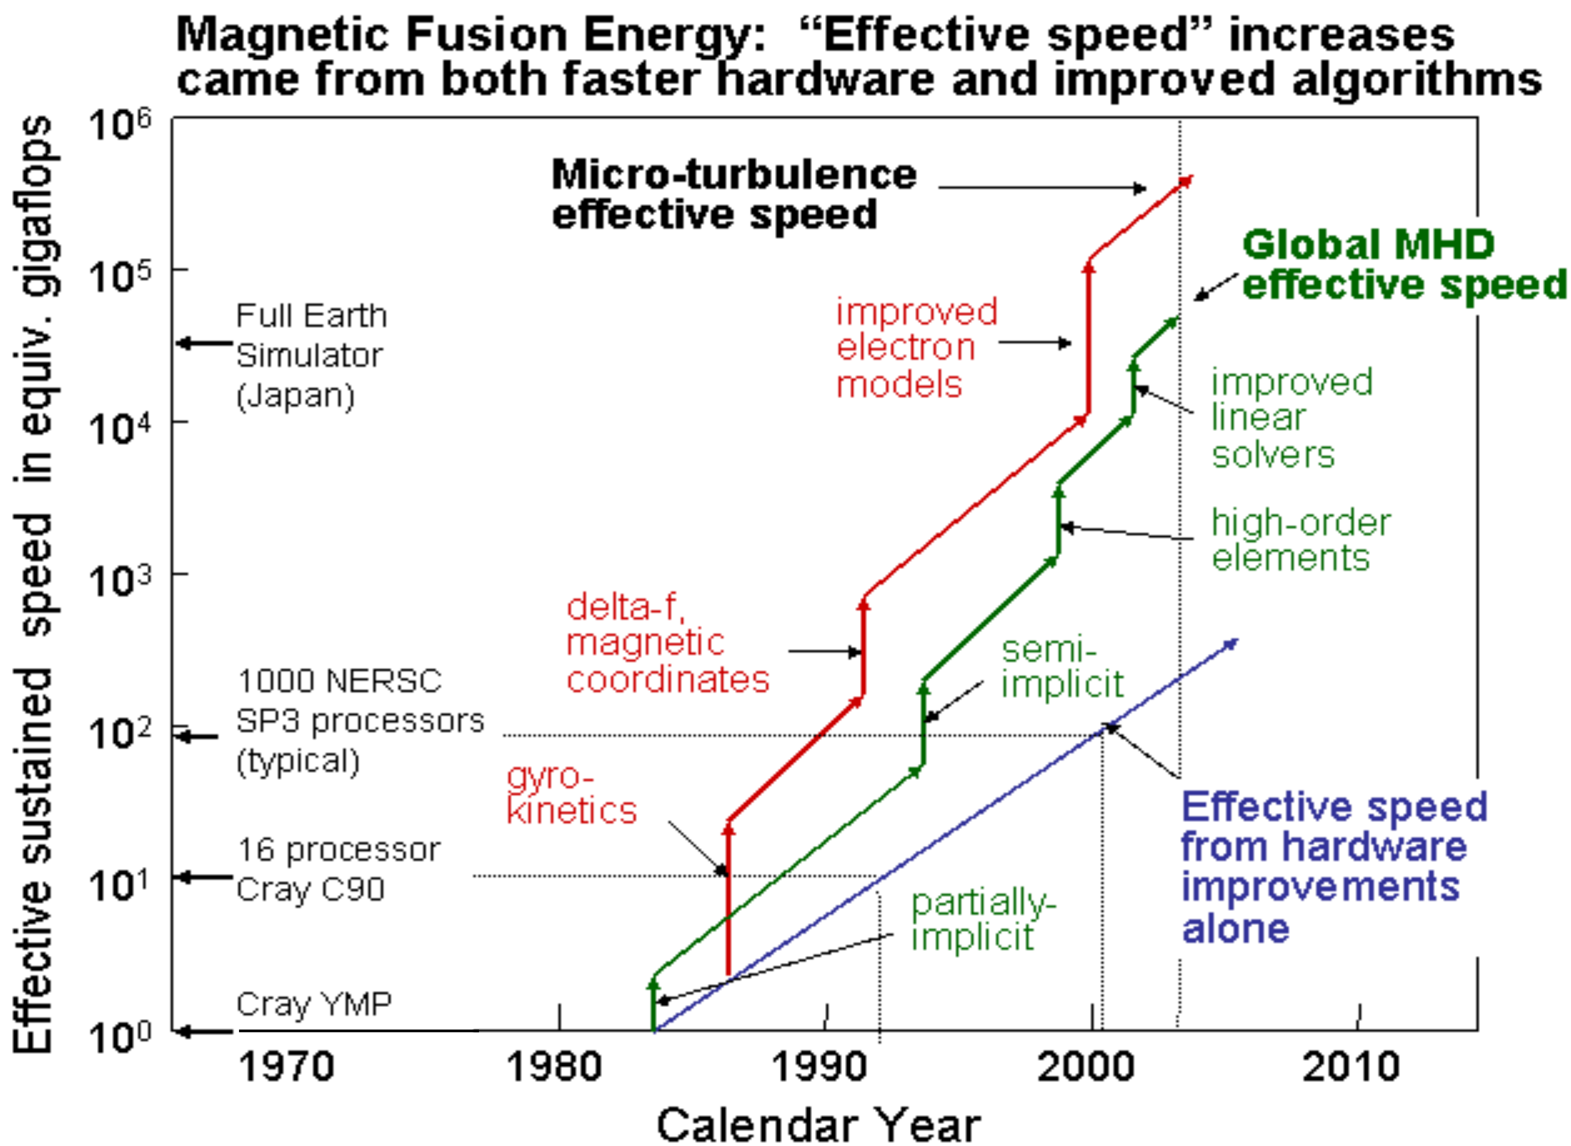

Figure from DOE "SCaLeS report" Volume 2 (Keyes et al., 2004)

"Moore's Law" for clean combustion simulations

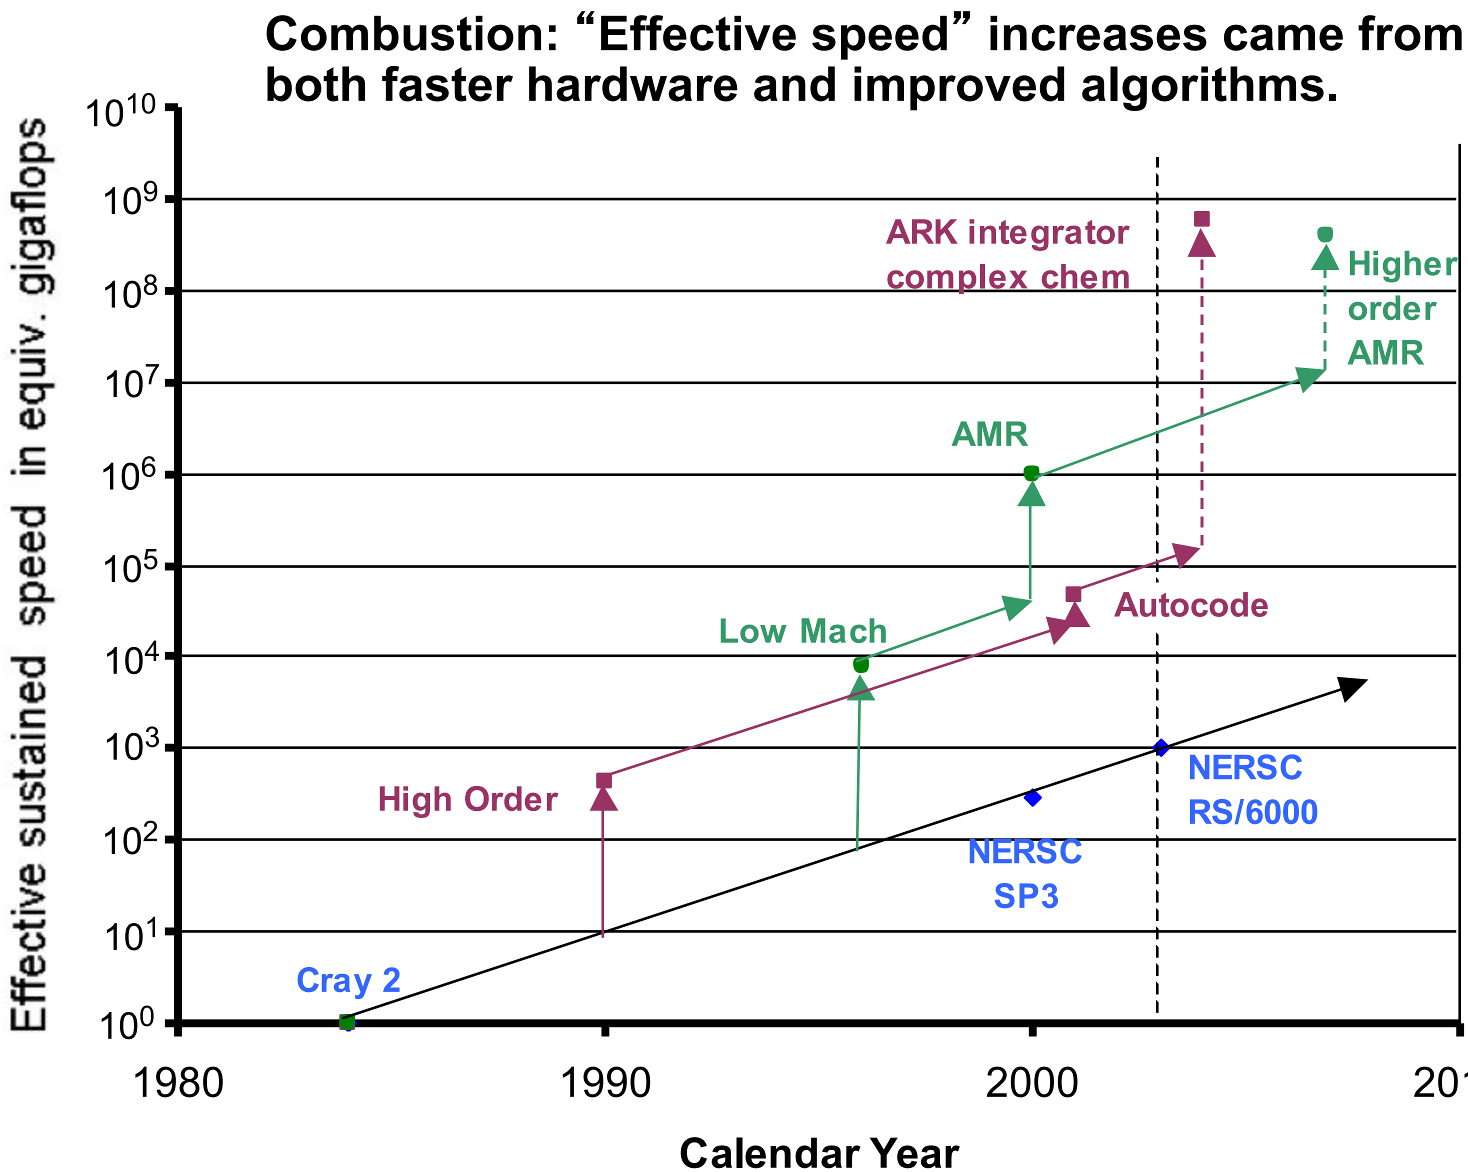

Figure from DOE "SCaLeS report" Volume 2 (Keyes et al., 2004)

Moore's Law and numerical algorithms

- First popularized in the 1992 NITRD bluebook: apply successive generations of algorithms to a fixed problem ("Poisson equation")
- In 24 "doubling times" (1.5 years) for Moore's Law for transistor density, better algorithms (software) contributed as much as better hardware
- $2^{24} \approx 16$ million $\Rightarrow$ 6 months of computing now takes 1 second on fixed hardware*
- *Two* factors of 16 million each if the best *algorithm* runs on the best *hardware*!

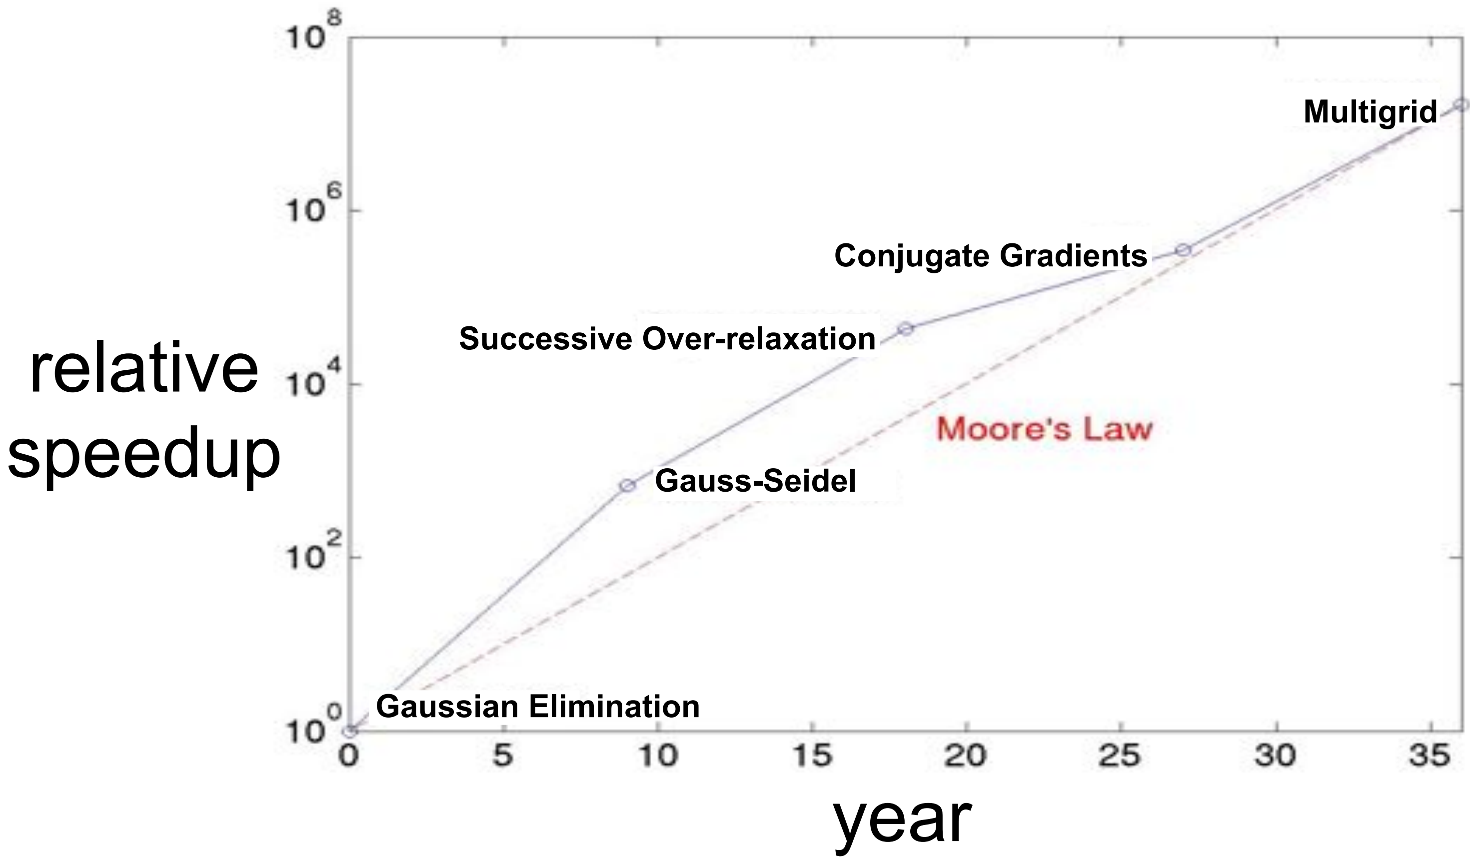

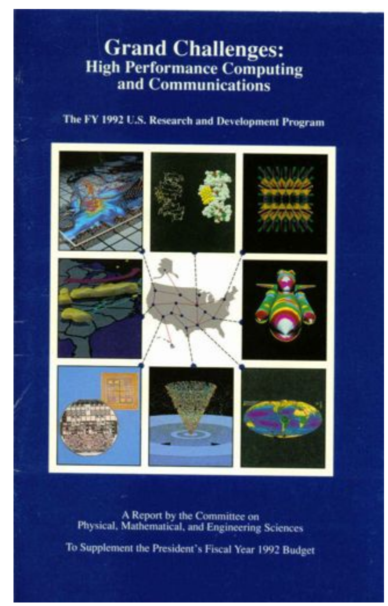

*algorithmic factor of improvement increases with problem size

These "Moore's Laws" get to the "bottom line"

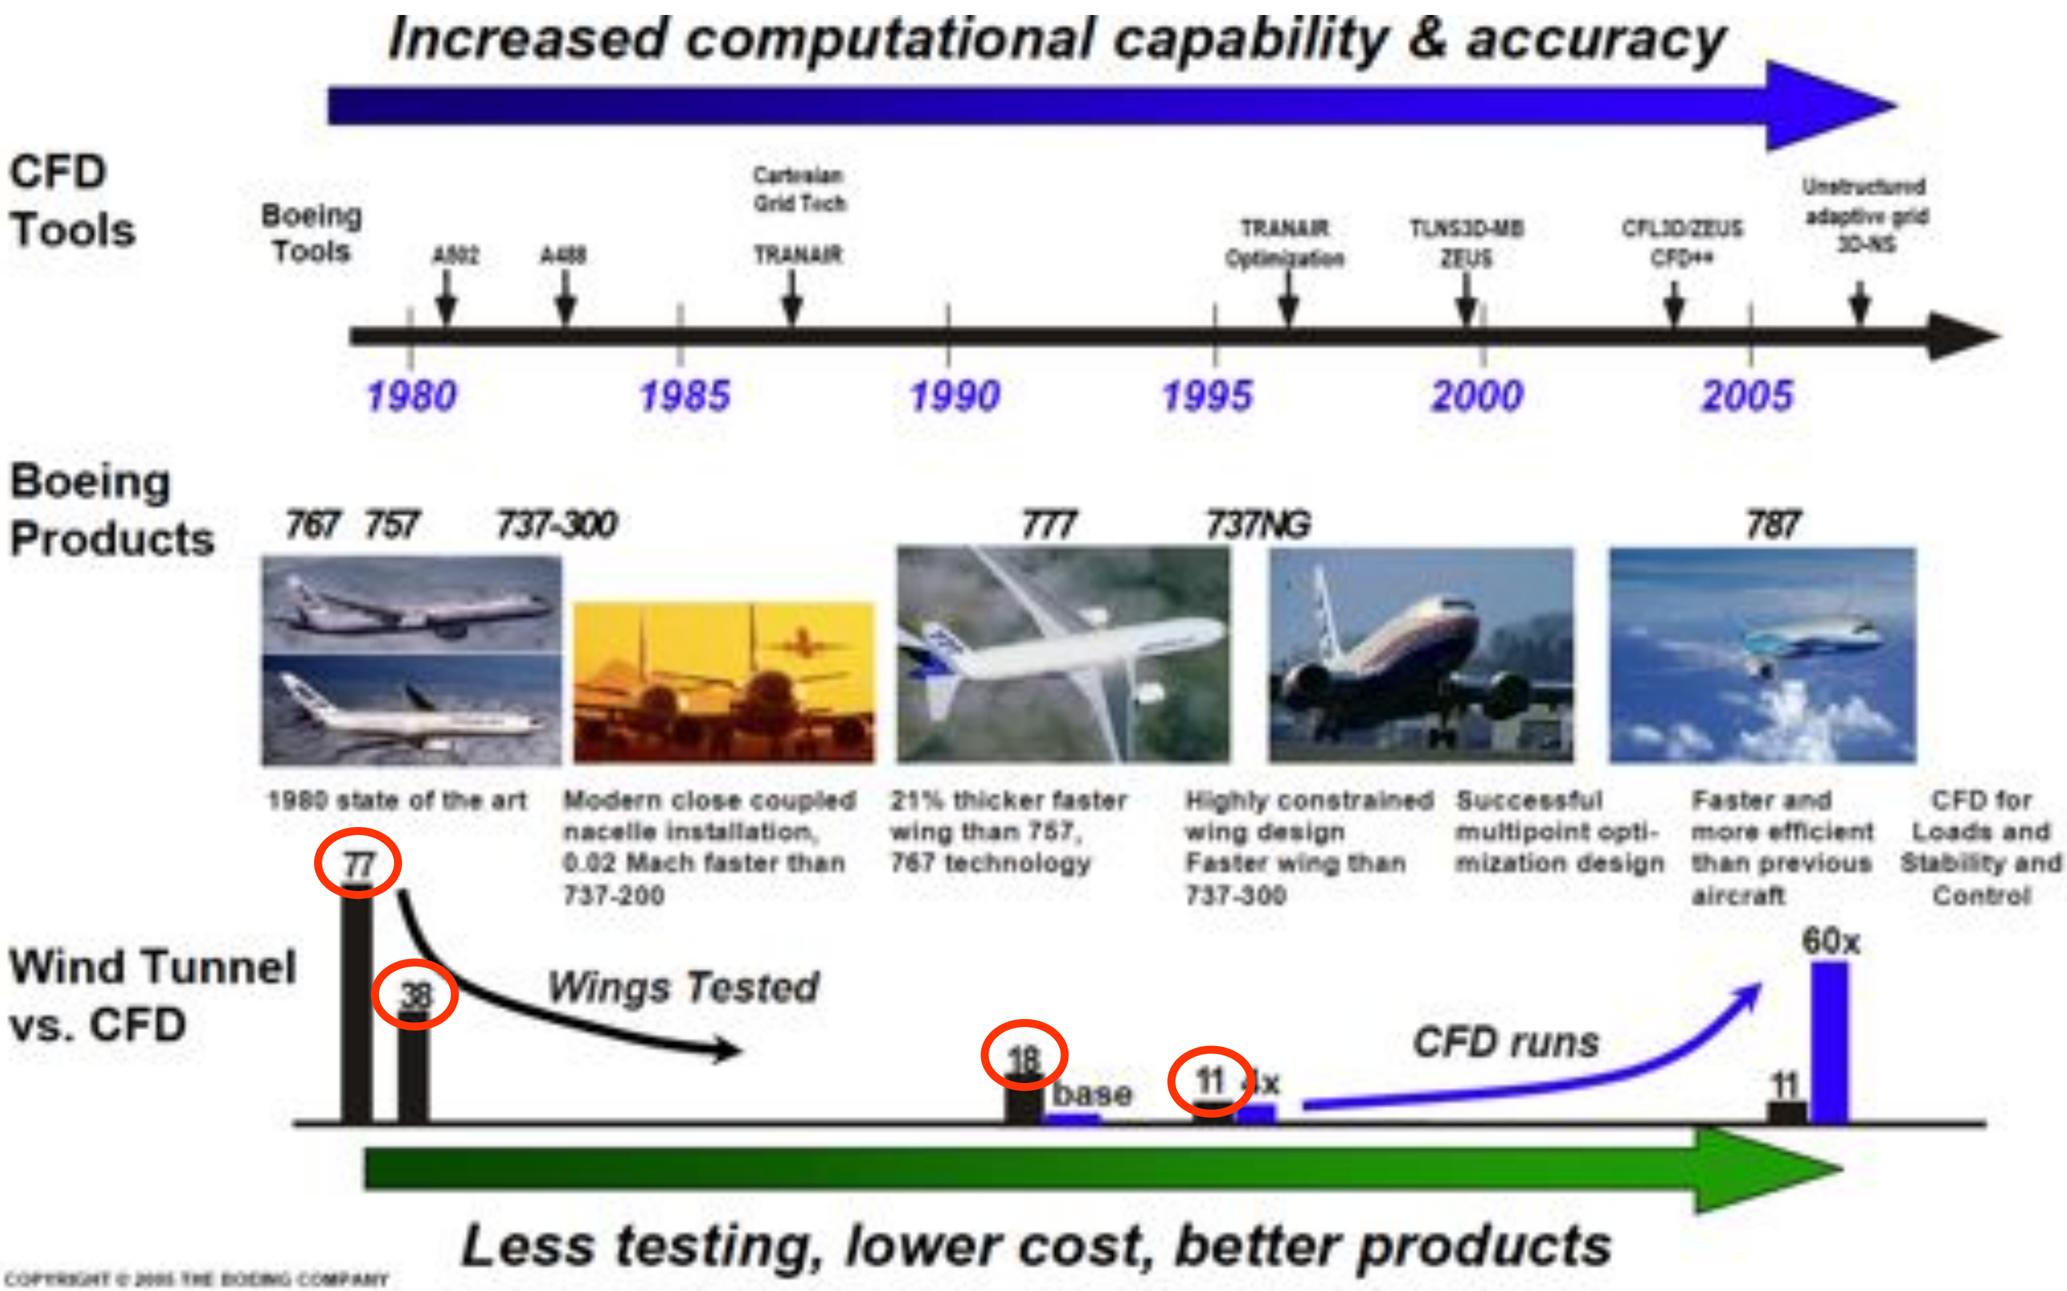

c/o Paul Johnson and Mark Goldhammer, Boeing

These "Moore's Laws" keep the country safe

Stockpile stewardship

Complex multi-material, unstructured electromagnetics

1.2 billion unknowns

2,000 processors

6 minutes – interactive timescale for engineers!

Not previously tractable at any turnaround

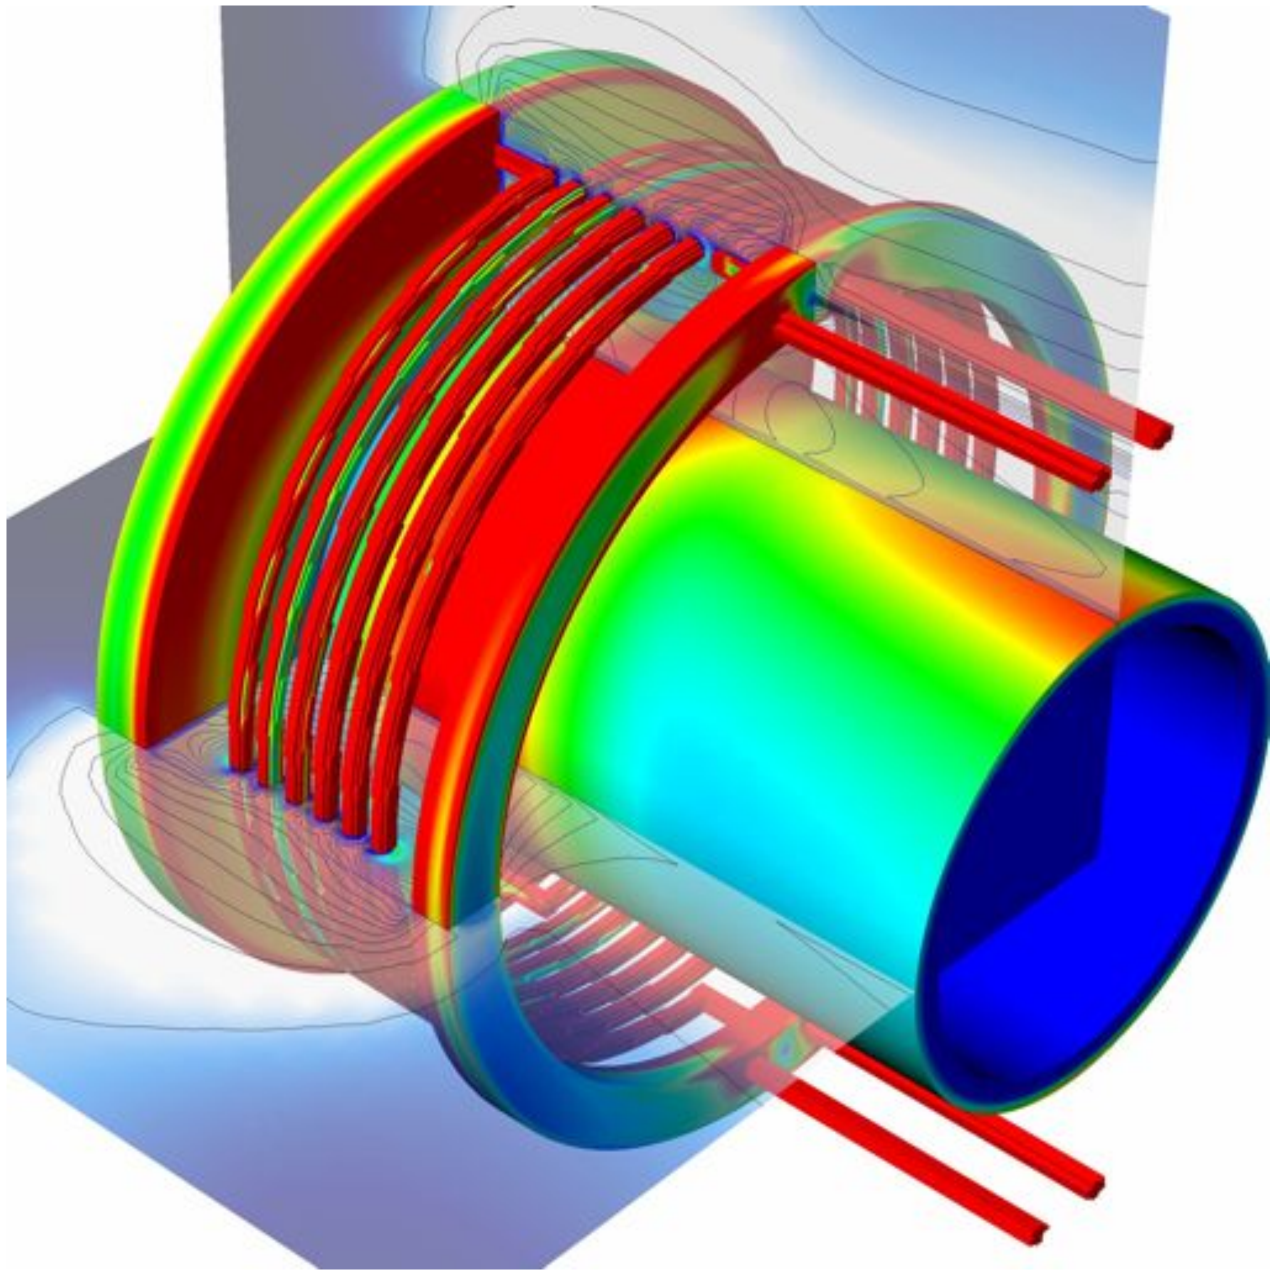

http://www.scidacreview.org/0901/pdf/bt.pdf

An exciting knowledge fusion for simulation

(dates are somewhat symbolic)

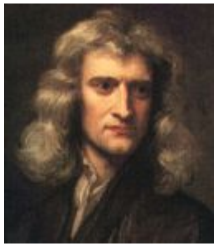

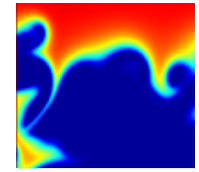

1686

scientific models

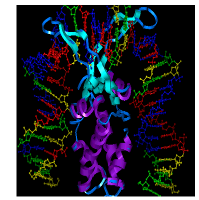

1947

numerical algorithms

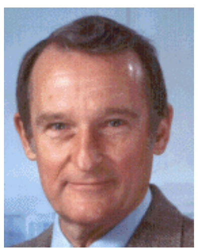

1976

computer architecture

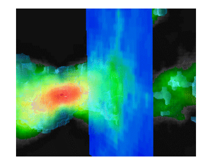

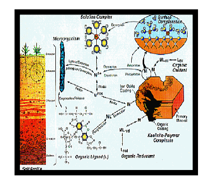

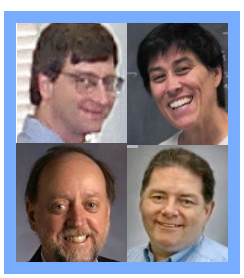

1992

scientific software engineering

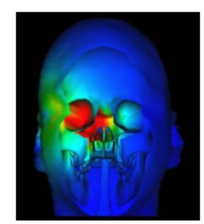

"Computational science is undergoing a phase transition."

Ecosystem: industry, national labs, academia

- **Mission-oriented and idea-oriented organizations make great partners for scientific discovery and technological advance**
- **No country does it better than the USA**

- **basic/applied**
- **short-term/long-term**
- **incubate/curate**
- **feedcorn/seedcorn**

In most countries, the barriers between basic and applied are much higher – even within academia!

Reward structures discourage exchanges, internships, cross-pollinations critical to innovation.

Basic research deposits into "treasury of ideas"

- **"Treasury" opened as scientists adapt to opportunities and constraints**
  - **driven by limiting resource (processing, storage, bandwidth, etc.), which is cyclical**
- **Algorithms arise to fill the gap**
  - **between *architectures that are available* and *applications that must be executed***
- **Many algorithms are *mined from the literature*, rather than *invented***
  - **underlining the importance of *basic research***
- **Many algorithmic advances are driven by particular physical problems *outside of the academic sandbox***
  - **underlining the importance of *applied research***

Algorithm	Born	Why?	Reborn	Why?
Conjugate gradients	*1952*	*direct solver*	*1970s*	*iterative solver*
Schwarz Alternating procedure	*1869*	*existence proof*	*1980s*	*parallel solver*
Space-filling curves	*1890*	*topological curiosity*	*1990s*	*memory mapping function*

NITRD agency success story: PETSc (1992-)

- The Portable Extensible Toolkit for Scientific Computing (PETSc)
  - used in thousands of scientific and engineering codes
  - software structure has inspired countless other library developers
- Suite of distributed data structures and routines for the scalable solution of large systems of equations
- Has won R&D 100 award, been part of multiple Gordon Bell prizes, Best Paper prizes; its developers won DOE's E. O. Lawrence award in 2011
- Funded by Argonne National Laboratory, DOE, and NSF

Acoustics, Aerodynamics, Air Pollution, Arterial Flow, Bone Fractures, Brain Surgery, Cancer Surgery, Cancer Hyperthermia, Carbon Sequestration, Cardiology, Cell Function, Combustion, Concrete, Corrosion, Data Mining, Dentistry, Earth Quakes, Economics, Fission, Fluid Dynamics, Fusion, Glaciers, Ground Water Flow, Hydrology, Linguistics, Mantle Convection, Magnetic Films, Material Science, Medical Imaging, Ocean Dynamics, Oil Recovery, PageRank, Polymer Injection Molding, Polymeric Membranes, Quantum Computing, Seismology, Semiconductors, Rockets, Relativity, …

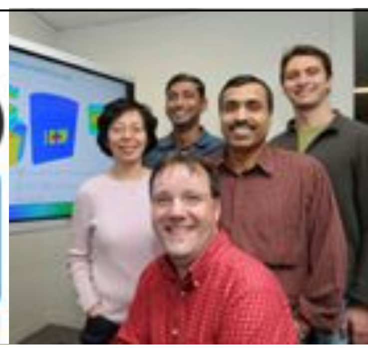

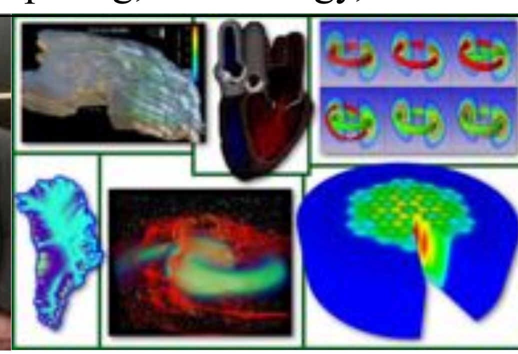

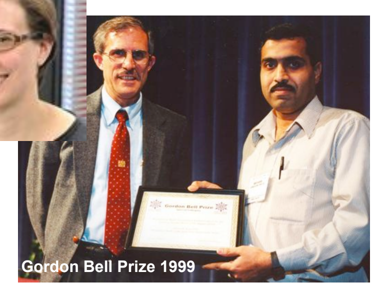

The "SciDAC" model of leveraging

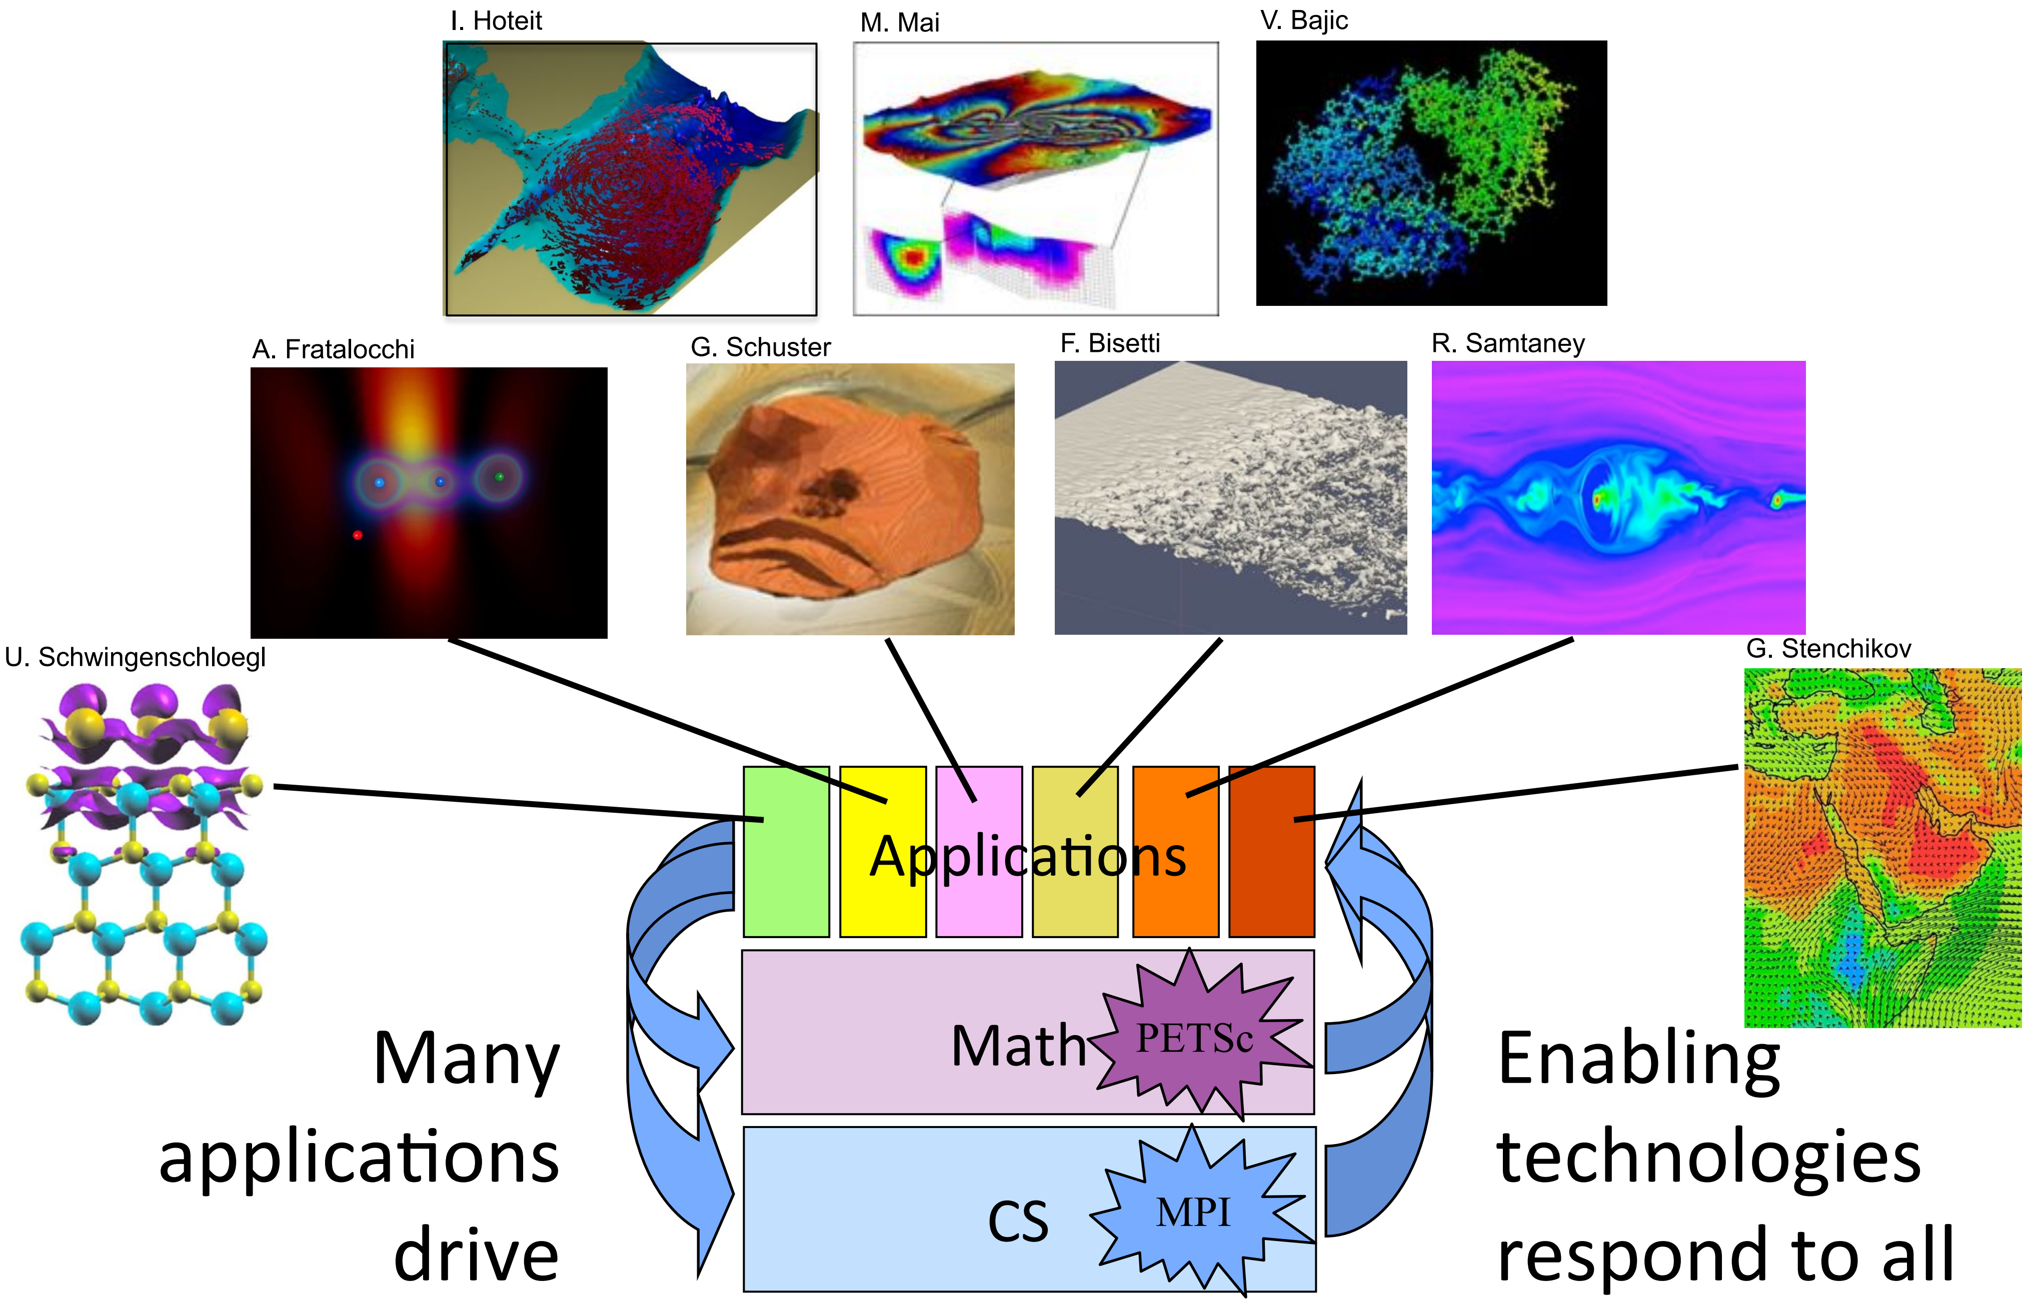

On to: the fourth pillar of scientific discovery

- **Data-enabled science**

"… Authors in this volume … refine an understanding of this new paradigm from a variety of disciplinary perspectives."
— *Gordon Bell, Microsoft Research*

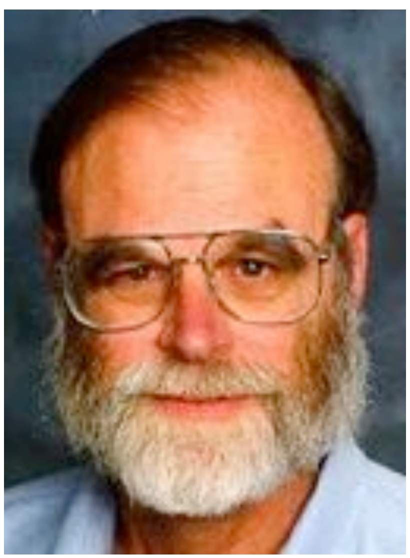

G. Bell, J. Gray, and A. Szalay, "Petascale Computational Systems: Balanced Cyber- Infrastructure in a Data-Centric World," in *IEEE Computer*, 39:110-112

Third and fourth paradigms belong together!

- **Future for simulation will embrace data culture**
  - **Inverse problems**
  - **Data assimilation**
  - **Uncertainty quantification**
  - **Immersive visualization and computational steering**
- **Future for data will embrace simulation culture**
  - **Simulation has mature culture of "optimal algorithms"**
  - **Complexity of solution or representation grows slowly in problem size or difficulty, *e.g., multigrid, multipole, FFT, sparse grids, spectral, interior point, solution-based adaptivity, importance sampling, etc.***
  - **Data analysis is at the beginning of optimality results**
  - **Complexity of solution or representation typically grows rapidly; but see recent breakthroughs, *e.g., wavelets, compressed sensing***
- **See Alex Szalay's talk for fourth paradigm**

The Tree and the Fruit

High performance computing is a phenomenally productive tree in the pursuit of scientific knowledge and technological advance.

High performance computing is also itself a fruit – an exciting fusion of computer science and mathematics that grows extraordinarily well in the USA.

NITRD rightfully enjoys the credit for envisioning and provisioning HPC for over twenty exhilarating years. May they stretch forward to the unimaginable.

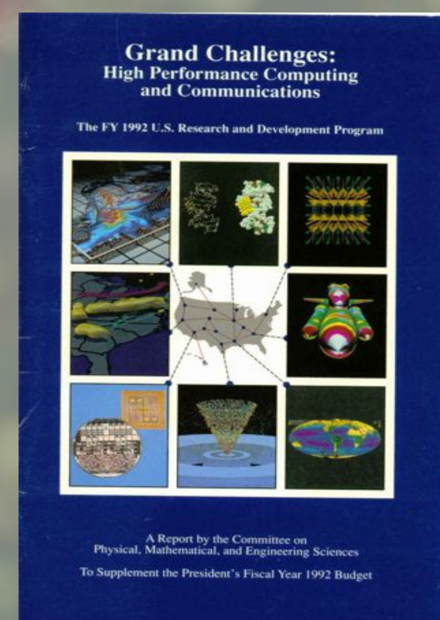

TRANSFORMING THE WORLD. DRIVING THE NATION'S COMPETITIVENESS. LEADING INTO THE FUTURE.

The Impact of NITRD

Two Decades of Game-Changing Breakthroughs in Networking and Information Technology — Expanding Possibilities Ahead